

ARGONNE NATIONAL LABORATORY
9700 South Cass Avenue
Argonne, Illinois 60439

**Development of a reduced-order design/optimization tool
for automotive engines using massively parallel computing**

Shashi M. Aithal¹ and Stefan M. Wild²

Mathematics and Computer Science Division
Preprint ANL/MCS-P5354-0615

June 2015

[to appear in Proceedings of the SAE 12th International Conference on
Engines & Vehicles (ICE2015), Capri, Italy]

<http://www.mcs.anl.gov/publications/>

²Mathematics and Computer Science Division, Argonne National Laboratory, Argonne, IL 60439, USA

¹Computing, Environmental, and Life Sciences Directorate, Argonne National Laboratory, Argonne, IL 60439, USA

Development of a Reduced-Order Design/Optimization Tool for Automotive Engines Using Massively Parallel Computing

Shashi M. Aithal and Stefan M. Wild

Argonne, National Laboratory, Argonne, IL 60439, USA

Copyright © 2015 SAE International

Abstract

Design and optimization of automotive engines present unique challenges on account of the large design space and conflicting constraints. A notable example of such a problem is optimizing the fuel consumption and reducing emissions over the drive cycle of an automotive engine. There are over twenty design variables (including operating conditions and geometry) for the above-mentioned problem. Conducting design, analyses, and optimization studies over such a large parametric space presents a serious computational challenge. The large design parameter space precludes the use of detailed numerical or experimental investigations. Physics-based reduced-order models can be used effectively in the design and optimization of such problems. Since a typical drive cycle is represented by 1500 to 2000 sample data points (engine cycles), it is essential to develop fast and robust computations so that the entire engine cycle computation is done close to real-time speeds (on the order of 100-150 milliseconds). Harnessing the power of high-performance computing, it is possible to perform optimization of automotive drive cycles using massively parallel computations. In this work, we discuss the development of a parallel fast and robust reduced-order modeling tool to compute integrated quantities such as fuel consumption and emissions (NO and CO) over a range of engine drive cycles. As an illustrative example, we perform a massively parallel simulation consisting of 4096 synthetic drive cycles, representative of a fleet of cars. The impact of parameters such as humidity, initial cylinder pressure, inlet air temperature, and residual gas fraction on the performance and emission are presented.

1. Introduction

Reduced-order models play an important role in many industrial applications, especially in the automotive industry. Reduced-order models are particularly important in situations where the design space is large since this size can preclude detailed numerical or experimental investigations. Design analyses and optimization of automotive engines present unique challenges on account of the large design space and conflicting constraints.

Optimizing the fuel consumption and reducing emissions over a drive cycle is a prototypical example of such a problem. Inlet pressure, equivalence ratio, humidity, EGR fraction (for Diesel engines), inlet air temperature, ignition timing, valve timing (which affects residual gas fraction), engine load, engine speed (RPM), along with engine dimensions (bore, stroke, compression ratio) constitute a vast parametric space for design/optimization studies. There are over twenty design variables and constraints that influence fuel economy and emissions. This large parametric space is further increased when one has to consider newer fuels and fuel blends (varying ratios of

fuel-additive mixtures) further complicating the design/optimization problem. Moreover, the effects of these parameters can vary significantly over a typical driving cycle, wherein the operating conditions (e.g., engine speed, load, injection/ignition timing) vary throughout the driving period. For instance, a typical 20-30-minute drive cycle could include a cold start for the engine, frequent starts and stops (traffic lights, intersections) and periods of constant speed/load (e.g., highway speeds). Such a drive cycle would experience a combination of engine speeds and loads (and hence fuel injected per engine cycle), humidity, and engine inlet air temperatures among other variables. Performing an optimization over the entire drive cycle to reduce emissions and fuel consumption for such a driving cycle is a daunting task. Since the engine speed varies typically from 800 to 3000 rpm (depending on the load), a 20-30-minute drive cycle would require about 20,000 to 40,000 engine cycles. For global analyses, engine speed and load (fuel consumption) may be recorded at intervals of 1 to 2 seconds, thus a typical drive cycle might contain 1000 to 1500 data points (of engine cycle data). For each sampled data point, which represents one engine cycle (compression and expansion), engine state variables (e.g., temperature, pressure, fuel-air mixture combination) need to be computed over 360 crank angle degrees (CAD), typically in intervals of roughly 0.5 CAD. These engine state variables are needed in order to compute the engine performance (e.g., torque and power) and engine-out emissions (e.g., NO and CO). Hence, each drive cycle requires the evaluation of over a million engine CAD (~1,500x720). Given the computational load, fast and robust reduced-order models describing the engine performance and emissions are required. In order to use such models for design analyses and optimization purposes, these drive cycle computations need to be completed in on the order of minutes, thus requiring each engine cycle computation to be completed in about 100-150 milliseconds. Furthermore, such a tool should be flexible, general-purpose, and user-friendly so as to enable its use for a wide range of operating conditions and a wide variety of fuels and fuel-additive combinations.

Several quasi-dimensional models have been developed since the early 1980s to study both gasoline and diesel engines (see, e.g., [1-11]) with varying degrees of fidelity. With few exceptions (e.g., [1,6,10]), however, these studies do not discuss the wall-clock time required for the computation of an engine cycle (a single compression and expansion stroke). The level of fidelity in the quasi-dimensional model and the wall-clock time required for the computation of an engine cycle are important considerations when one is interested in the simulation of a large number of drive cycles for design analyses and optimization, and real-time control. Furthermore, most 0-D, 1-D, and other quasi-dimensional codes described in the open literature are primarily aimed at solving a single engine cycle (–i.e., 360 CAD) for a given set of operating conditions.

The primary focus of this work was to develop a flexible, general-purpose, and user-friendly tool to simulate long drive cycles of vehicles operating under a wide range of operating conditions and fuel/fuel-additive combinations. Emphasis was laid on using fast, robust computations in order to complete an entire engine cycle computation in close to real-time speeds (on the order of 100-150 milliseconds). Since the design and state spaces are large, one has to harness the power of large-scale computing to reduce the overall wall-clock time associated with such computations. The availability of large computer clusters, and more recently, cloud computing, has made possible concurrent evaluation of a large combination of parametric cases and streamlined analyses of the results in a matter of minutes, rather than days or months. There are relatively few studies that use large-scale computing for system-level analyses [12-13]. To the knowledge of the authors, this work is the first to implement and demonstrate the use of physics-based engine models for large-scale analysis of a fleet of cars using concurrent computations of the various parametric cases.

This work describes the development of pMODES (**parallel Multi-Fuel Otto Diesel Engine Simulator**), a framework aimed at accomplishing the goals mentioned above. This work builds on the modeling tool described in [1]. Ref. [1] describes the various sub-models used in pMODES along with detailed validations. All the sub-models described in Ref. [1] are widely used in other quasi-dimensional models and their accuracy and reliability has been well-documented. Furthermore, the modular architecture of pMODES enables users to implement other sub-models of their choice. pMODES includes several enhancements to the tool described in [1] in order to enable its use for a larger spectrum of design, analysis, optimization, and real-time control studies. The design tool has been fully parallelized in order to streamline the workflow management to concurrently conduct large-scale parametric sweeps of the drive cycle of a fleet of cars. The set of species included in the combusting mixture has been expanded to include charged species. A set of seven charged species that enables the computation of the temporal variation of ion current and location (in CAD) of the ion-current peak for engine diagnostics and control has been implemented. The set of neutral species considered in the computations includes species such as C, CH, and C₂H₂, which allows the inclusion of simplified soot models as well.

The above-mentioned features in pMODES enable massively parallel simulations of the drive cycles of a fleet of cars with varying operating conditions. pMODES can thus be used as a virtual dyno, which can complement experimental dyno data or be used to conduct large parametric sweeps over a range of engine design parameters with minimal user intervention. pMODES can also be used for detailed design analyses for next-generation engines. As an illustrative example, we present results obtained by conducting a large-scale parametric sweep over synthetic drive cycles for a single-cylinder gasoline engine (multi-cylinder engines can also be studied with this tool). Four different values of each of the four variables, namely, initial cylinder pressure, humidity ratio, inlet air temperature, and residual gas fraction (RGF) were varied for sixteen different synthetic drive cycles giving rise to 4,096 cases in all (4⁴ × 16 = 4,096). These cases were run on 4,096 processors on a mid-sized cluster with an Intel Sandy Bridge architecture and on the IBM Blue Gene/Q (BG/Q) supercomputer with PowerPC A2 1600 MHz processors, both at Argonne National Laboratory.

This paper is organized as follows. Sections 2 and 3 briefly describe the model and the method of solution. Section 4 describes

sample results for the single-cylinder gasoline engine followed by a summary and conclusion of the main findings of this work.

2. Model Description

Figure 1 shows the solution methodology of pMODES. As seen in Figure 1, the main physics-based engine module computes the temporal variation of the important engine state variables, namely, temperature (burned and unburned gas), pressure, and mixture composition. This information is used to compute emissions (e.g., NO, CO, soot) and ion current using reduced-order models. Optionally, the engine wall temperature can also be computed as described in [2]. Furthermore, the computed cylinder pressure can be used to evaluate engine performance metrics such as BMEP, torque, and power. Details of the solution procedure for obtaining the cylinder state variables along with emissions, wall temperature and ion current (or voltage) and are described in [1, 2, 14, 15]. Detailed validation of the ion-current sub-models used in this work have been discussed in [14, 15] and hence will not be presented here. Briefly, the energy equation describing the relationship between the crank angle and cylinder pressure is solved to obtain the temporal variation of engine pressure along with the burned and unburned gas temperature as shown in Eqns (1) – (3).

$$\frac{dP(\theta)}{d\theta} = \frac{\gamma - 1}{V(\theta)} (Q_{in} - Q_{loss}) - \gamma \frac{P(\theta)}{V(\theta)} \frac{dV}{d\theta} \quad (1)$$

$$T_b(\theta) = \frac{P(\theta) V_b(\theta)}{m_b(\theta) R_{gb}} \quad (2)$$

$$T_u(\theta) = \frac{P(\theta) V_u(\theta)}{m_u(\theta) R_{gu}} \quad (3)$$

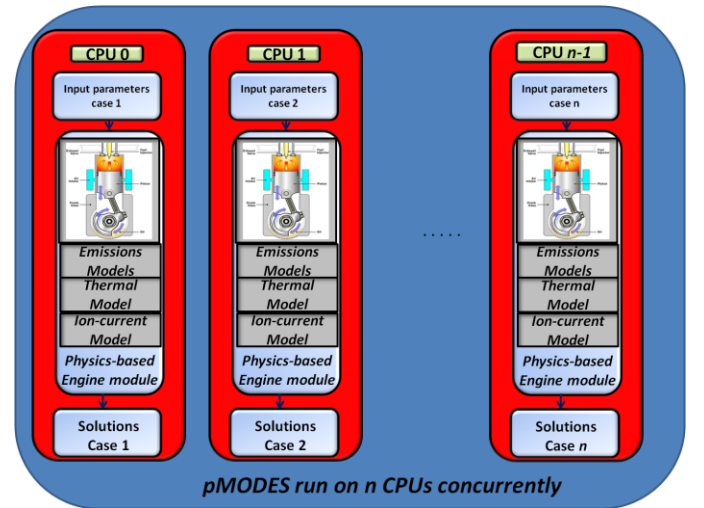


Figure 1: Solution methodology of pMODES.

Knowing the initial cylinder mixture composition (air-fuel ratio, EGR fraction, humidity ratio, and residual gas fraction) and the fuel burn rate (as described by a user-defined Weibe function), it is possible to compute the temporal variation of the composition of the burned gas. Based on the mixture composition, temperature, and pressure, it is possible to compute the equilibrium composition of the burned gas at a given crank angle. The equilibrium composition is

used to compute emissions and ion currents based on reduced-order models as described Eqns (4) - (6).

$$\begin{aligned} \frac{d[NO]}{dt} = & k_1^+ [O][N_2] + k_2^+ [N][O_2] + k_3^+ [N][OH] \\ & - k_1^- [NO][N] - k_2^- [NO][O] - k_3^- [NO][H] \\ & - [NO] \frac{1}{V_b} \frac{dV_b}{dt} \end{aligned} \quad (4)$$

$$\begin{aligned} \frac{d[CO]}{dt} \Big|_{cv} = & (C_{\omega} R_1 + R_2 + R_3) \left(1 - \frac{[CO]}{[CO]_{eq}} \right) \\ & - [CO] \frac{1}{V_b} \frac{dV_b}{dt} \end{aligned} \quad (5)$$

In the above equation, $R_1 = (k_{f,1a} + k_{f,1b} + k_{f,1c})[CO]_{eq}[OH]_{eq}$, $R_2 = k_{b,2}[CO]_{eq}[O_2]_{eq}$, $R_3 = k_{f,3}[CO]_{eq}[O]_{eq}[M]_{eq}$, and $[M] = [H_2]_{eq} + 6.5[H_2O]_{eq} + 0.4[O_2]_{eq} + 0.4[N_2]_{eq}$. All the concentrations (such as $[N_2]_{eq}$) are equilibrium values at a given temperature and pressure corresponding to a given crank angle.

$$I_s(t) = K \left(\sum_{j=1}^{N_C} e_j n_j \mu_j \right) \quad (6)$$

3. Method of Solution

The main solver in pMODES has been designed to be highly modular in order to ensure portability and extensibility. Special attention was paid to the code layout and dataflow to ensure that new models for emissions and ion current can be included at a later stage with minimal changes in the overall structure of the code. As shown in Figure 1, each parametric case of pMODES is run as a unique MPI rank (in our study, the number of ranks is set to the number of cores). This enables a large number of cases to be run concurrently for design analyses and optimization studies. Each MPI rank reads its input parameters (design variables such as initial cylinder pressure, air temperature, humidity ratio, etc.) along with the fuel input per cycle, engine speed, ignition/injection timing, and valve timing (drive-cycle information), from a uniquely named input file, completes the computation, and writes the solution to a unique solution file. This I/O design was adopted to minimize data communication and also for ease of post-processing of the computed variables. Robust solution techniques were used to reduce the overall computational time and optimize the single-processor performance. The number of parametric cases to be conducted concurrently is decided by the user at runtime by assigning 'n', the number of processors. This makes pMODES highly flexible and hence suitable for various computer architectures, from desktops to leadership-class machines (supercomputers) equipped with many thousands of cores. The simple and self-contained code structure makes pMODES ideally suited for cloud computing. Over 90% of the computational time is spent in computing the equilibrium concentration of the twenty six species assumed to be present in the combustion mixture. A modified Newton-Raphson scheme with a robust, physics-based relaxation methodology was used to accelerate convergence and reduce the computational time required to obtain the equilibrium composition of the mixture. The system of equations, cast in a matrix-vector form as discussed in [15] was solved using an

optimized LAPACK subroutine, thus greatly increasing the portability of pMODES to various architectures.

4. Results and Discussions

This section describes the use of pMODES to study the performance and emission characteristics of a fleet of cars consisting of 4096 different synthetic drive cycles. A single-cylinder gasoline engine with fixed dimensions is used for this illustrative example. Spark timing was fixed at 26 bTDC and the combustion parameters, such as combustion duration and Wiebe parameters, were kept constant for all the cases considered. Spark timing, combustion parameters, and engine dimensions are all design variables that can be user-defined parameters but which have been kept constant for this study for the sake of simplicity. The impact of varying four parameters (namely, humidity, initial cylinder pressure, inlet air temperature and residual gas fraction) on the performance, emissions, and ion current are presented. The combination of these four parameters was chosen for illustrative purposes only. The initial cylinder gas composition (moles of O_2 , N_2 , H_2O , CO_2 , and fuel) depend on the choice of these parameters. This particular case-study thus demonstrates the impact of initial cylinder composition on the engine performance and emissions. pMODES can also be used to study the parametric variation(s) of any other arbitrary set of engine variable(s) based on the user's requirements or engine experimental data. For instance, instead of fixing the initial cylinder pressure, one could conduct a parametric simulation of the air-flow into the cylinder. In such a case, pMODES would calculate the initial cylinder pressure based on the inlet air temperature, residual gas fraction and humidity. Table 1 shows the parameters and their range of variation. Each of these drive cycles has 1500 data points with each data point representing the fuel injected per cycle. Thus, if it is assumed that the data is sampled every second, the drive cycle with 1500 data points corresponds to 1500 seconds (25 minutes) of engine test time. The same procedure can be used to study multi-cylinder engines.

Table 1. List of parameters and their range of variation.

Variables				
Case	Initial cylinder pressure, p (atm)	Humidity ratio, w (%)	Inlet air temperature, T_c (C)	RGF (%)
1	0.88	0.0	28	0
2	0.92	1.0	29	1
3	0.95	2.0	30	2
4	1.0	2.75	31	3

A wealth of information about the impact of varying these parameters independently over the operating range on emissions and performance can be obtained. We discuss cases wherein a single synthetic drive cycle is subjected to a variation of each of the parameters shown in Table 1 while keeping the others constant. The impact of each parameter varied over its prescribed range on the engine-out NO and CO is described next.

Figure 2 shows the variation of the fuel input (in mg/s) in a synthetic drive cycle and the corresponding effect of varying the humidity ratio on the engine-out NO and CO. The humidity ratio (w) is varied from 0% (dry air) to 2.75% (near saturated air). The initial

cylinder pressure was maintained at 0.88 atm, temperature was fixed at 28°C and the RGF was set to zero. Figure 2 shows that increasing the humidity ratio decreases the engine-out NO while it increases the engine-out CO. Figure 2 shows that the reduction in the engine-out NO (from dry air to saturated air) could be as high as 20-25%.

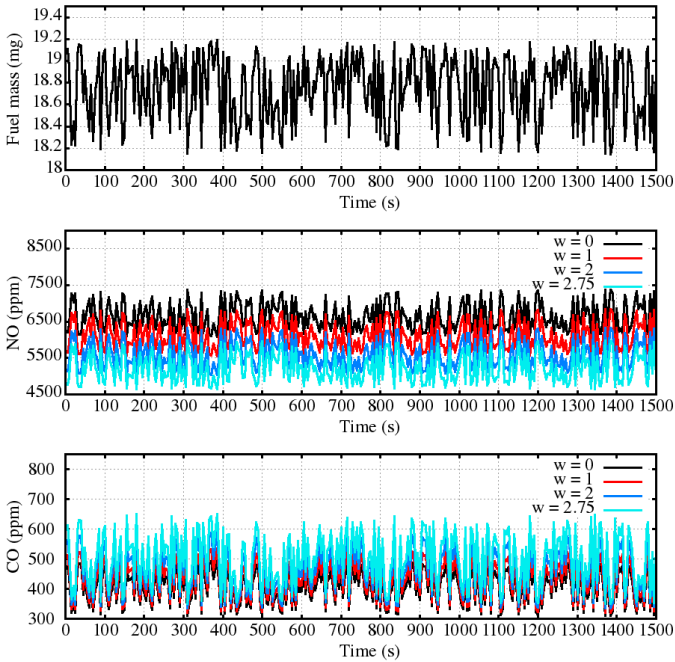


Figure 2: Effect of humidity on engine-out NO and CO.

The reduction in NO with increase in humidity is expected. Since the initial cylinder pressure and air inlet temperature are fixed, increasing the fraction of humidity (H_2O) reduces the mass of dry air and hence the N_2 and O_2 content in the cylinder. Thus, increased humidity has an effect similar to the introduction of EGR, leading to reduced engine-out NO. This information can be used by control engineers to adjust the required EGR fraction for a given target level of NO reduction, thus minimizing the harmful effects of EGR on the engine. Figure 2 also shows that increasing the humidity ratio increases the CO by as much as 20-25%. The increase in engine-out CO can be explained on the basis of the reduced-order CO model described in Eqn. 5. It is seen that the term R_3 has a strong dependence on the concentration of H_2O (via the term $[M]$) and is directly proportional to the production rate of CO as shown in Eqn. (5). The higher humidity level implies an increase in R_3 leading to increased levels of engine-out CO. Integration of the area under the NO and CO curves would yield the overall engine-out NO and CO in mg (with appropriate conversions from PPM to mg) for a given drive cycle or a range of drive cycles. This information can be used by design engineers to forecast emissions while adhering to constraints imposed by CAFE (Corporate Average Fuel Economy) regulations for newer fleet of cars.

For the sake of clarity, Figure 3 shows a shorter snapshot of Figure 2 over a period of 50 seconds. The impact of humidity ratio on the NO and CO is seen more clearly in Figure 3.

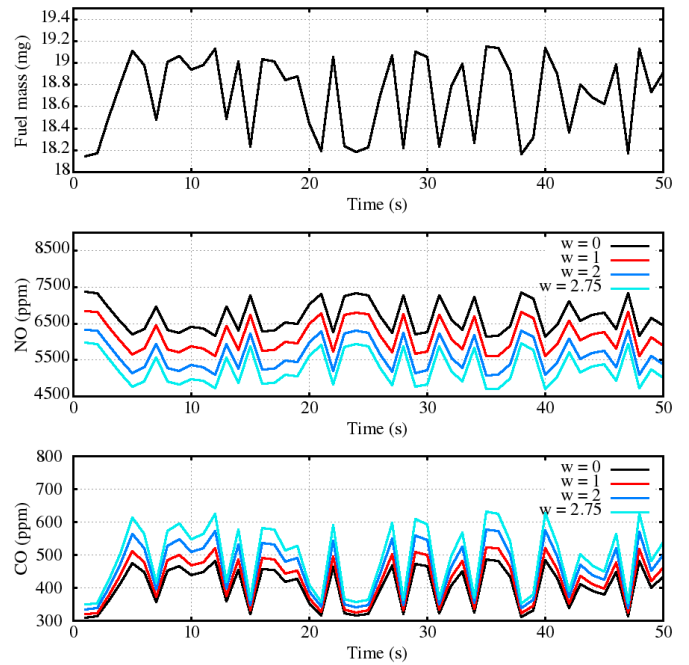


Figure 3: Magnified view (shorter time scale) of Figure 2.

Figure 4 shows the effect of varying the initial cylinder pressure on the engine-out NO and CO for a fixed humidity ratio ($w = 0\%$), inlet air temperature (28 °C), and RGF (0%). It is seen that increasing the initial cylinder pressure, for the same fuel input, implies greater initial air mass (on account of the higher fixed initial cylinder pressure). As the initial cylinder pressure is increased, the concentrations of N_2 and O_2 increase, thus leading to higher values of engine-out NO. It is also seen that the lower the initial cylinder pressure, the higher are the values of engine-out CO.

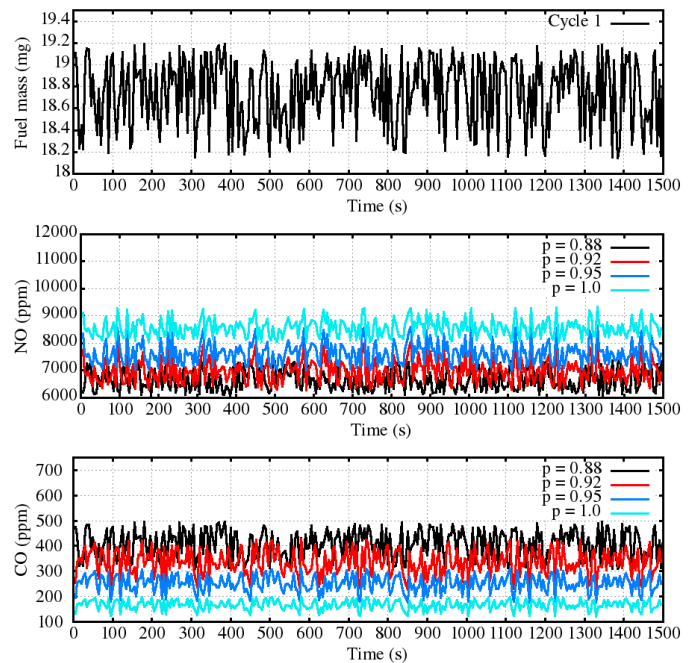


Figure 4: Effect of initial cylinder pressure (p) on NO and CO.

Figure 5 shows the effect of inlet air temperature (T_c) on NO and CO for the same drive cycle with $p = 0.88$, $w = 0\%$, and RGF

= 0%. It is seen that variation of the inlet air temperature has a minimal effect on the engine-out NO and CO. From Figures 2, 4, and

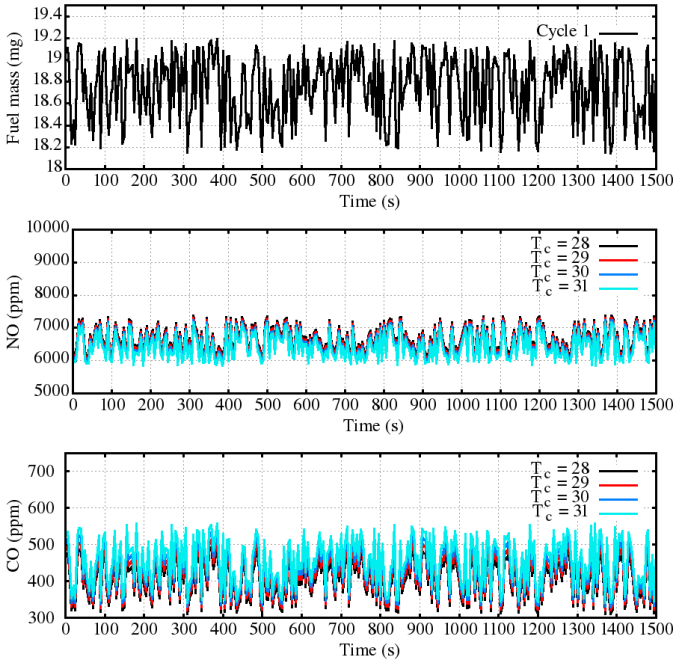


Figure 5: Effect of inlet air temperature (T_c) on NO and CO.

5 it is seen that it is possible to obtain good qualitative information about the impact of various operating conditions on the engine-out emissions. With well-calibrated models, detailed quantitative information can also be obtained from such studies. From Figures 2 to 5, it is seen that minimizing NO and CO can lead to conflicting constraints (a design parameter reducing NO leads to increasing CO). The large-scale parametric simulation of long drive cycles enables design engineers to evaluate an optimum operating envelope to meet regulatory constraints on fuel consumption and emissions while optimizing engine performance.

In addition to the impact of various operating parameters on emissions, it is also possible to ascertain parameters such as exhaust temperature and pressure corresponding at EVO (exhaust valve opening). The temperature and pressure at EVO directly impact engine efficiency and hence an important design consideration. Figure 6 shows the impact of humidity on the exhaust temperature (only 20 sec has been shown for clarity). As it can be seen from Figure 6, dry air ($w = 0\%$) leads to a lower exhaust temperature as compared to near saturated air. Since the cylinder pressure is kept constant at 0.88 atm for both the cases (dry and saturated air), the initial mass of air in the cylinder is higher for dry air as compared to near saturated humid air. The increased air mass leads to a low exhaust temperature. Since exhaust temperatures impact the overall engine performance, it might be important to consider their role in the overall optimization problem (engine performance and emissions).




Figure 6: Effect of humidity on the exhaust temperature.

In addition to performance and emission characteristics, pMODES can be used for engine diagnostics and control applications. Ion current (or voltage) signals have been used to evaluate the combustion quality and control in engines [16-19]. pMODES computes the temporal variation of important charged species required to compute the ion current signal in engines. The peak ion signal and its location computed by pMODES can be used for formulating a control strategy. It can also be used for advanced engine diagnostics by correlating the location and magnitude of the current peak with other engine state variables, such as peak pressure, temperature, ignition timing, and equivalence ratio. Figure 7 shows the impact of RGF on the ion current for a fixed drive cycle with the inlet temperature set at 28C, relative humidity set at 0%, and the initial cylinder pressure set at 0.88 atm. It is seen that increasing the RGF from 0% to 3%, yields a higher current signal. This trend can be explained by noting that presence of RGF reduces the fraction of molecular oxygen (O_2) in the cylinder for a given inlet cylinder pressure. The magnitude of the current signal is strongly influenced by the electron concentration in the mixture on account of their higher mobility. However, the concentration of electrons in a combustion mixture is strongly influenced by the presence of O_2 . Electrons in a combustion mixture attach themselves to O_2 via three-body collisions ($e^- + M + O_2 \rightarrow O_2^- + M$). Thus, a higher concentration of O_2 in the mixture (as in the case with no RGF) leads to increased electron attachment, leading to a lower concentration of electrons and hence a lower magnitude of ion current. Ion current traces such as those shown in Figure 7 can thus be used for diagnostic and control purposes.




Figure 7: Effect of RGF on ion current.

Thus far we have discussed the impact of a single design parameter (e.g., humidity, initial cylinder pressure, RGF) on the performance, emission, and diagnostic characteristics. More complex correlations based on the simultaneous variation of two or more design parameters can also be studied. In addition to understanding and evaluating the impact of design parameter(s) on the optimization objectives (fuel economy and emissions), a wealth of information can be obtained from a large parametric sweep over different drive cycles. Designers can use “filtering” to extract basic statistics and identify input configurations of interest. For instance, one can determine peak temperatures and pressures over an operating envelope and ensure that the temperature peaks do not exceed prescribed design limits. The filtering process can be used in the design phase of newer engines to isolate operating conditions that yield peaks above a threshold. Empirical distribution characterization can be used to extract cross-configuration information. Such distributions can be used, for example, to determine fleet-wide fuel economy [20] or to characterize emissions as a function of ambient pressure. Furthermore, one can conduct sensitivity analysis to determine how operating conditions or other independent variables affect observables of interest. Sensitivity analysis can be used, for example, to determine the impact of different fuels and fuel-additive combinations on the performance and emission characteristics of a fleet of cars. Using the data generated from these large parametric sweeps, tradeoff visualization and analysis can be used to identify a configuration that has worse figures of merit as compared to some other configuration. Such analyses can also be used, for example, to identify vehicle configurations wherein a small penalty in performance would yield substantial gains in fuel economy or reduction in emissions.

The capabilities of pMODES can be used as one “closes the loop” between the simulation and analysis for purposes of simulation-based design optimization [2] or optimal experimental design to determine configurations that should be tested on a dyno. Furthermore, pMODES can be used to generate input configurations in order to optimize a design objective of interest. Distributional information can be used to generate scenarios (e.g., ambient or operating conditions,

drive cycle variations) for use in sample average approximation for optimization under uncertainty [21]. Similarly, tradeoff analysis forms the basis for simultaneously optimizing multiple conflicting objectives [22,23], such as performance and engine lifetime/reliability.

As pointed out earlier, computational speed, scalability, and portability are important considerations in the design of pMODES. Indeed, these attributes are important when conducting large-scale parametric sweeps for design analyses and optimization studies. In order to test the computational speed, scalability, and portability, pMODES was run on two different machines, namely, a mid-sized cluster with ~ 5000 Intel Sandy Bridge 2.6GHz cores, and supercomputer (IBM, Blue Gene Q) with ~ 786000 cores powered by PowerPC A2 1.6 GHz processors. In addition to testing the portability, running these simulations on two different machines yields important performance metrics (e.g., wall-time) for different architectures and compilers. Up to 4096 different parametric cases were run concurrently on each of these machines. Table 2 shows the total wall time for evaluating a number of drive cycles equal to the number of cores (with 1500 data points) on each of these systems. It is seen that for both machines considered, as the number of cores were increased with a proportional increase in the number of cases, the total wall-time for the concurrent computation was nearly constant. This is to be expected since each of the parametric cases is independent of the others. Table 2 thus shows excellent weak scaling on both machines. We attribute the slight increase in overall computational time as the number of cores (and thus cases) increases primarily to imbalances in individual case solution times and to increased contention for the I/O operations. Furthermore, we see greater imbalances across cases for the mid-sized cluster as compared to the supercomputer.

Table 2. Timing studies on IBM (BG/Q) and Intel Sandy Bridge cores.

Cores	Sandy Bridge (min:sec)	IBM (BG/Q) (min:sec)
32	3:15	-
64	3:20	-
128	3:20	-
256	3:20	-
512	3:32	18:14
1024	3:46	18:52
2048	3:53	19:01
4096	4:26	19:08

Table 2 shows that the time to simulate an individual drive cycle on the IBM BG/Q is about 6 times longer than the Intel Sandy Bridge. This difference can be attributed to the hardware and software differences between the two machines (clock-speed, memory, compilers, etc.). However, it should be noted that the computation of the entire drive cycle estimated to be about 25 minutes of real-time is accomplished in less than 20 minutes with each of the machines, thus demonstrating near real-time computation speeds. The supercomputer, with nearly 786000 cores, would enable a larger set of concurrent calculations for more massive parametric studies as compared to the mid-sized cluster. Since massive parametric sweeps can be conducted at near real-time speeds, this

work demonstrates an important step forward in the use of massively parallel computing for design analyses and optimization of automotive engines.

5. Summary and Conclusions

This work discussed the development of pMODES, a parallel, reduced-order modeling engine modeling tool. pMODES is primarily aimed at concurrently conducting large-scale parametric studies for design analyses and optimization studies involving a fleet of automotive engines. pMODES has a modular structure that enables its extension to new reduced-order models for computing emissions and performance of engines. Robust solvers enable fast and accurate computation of the equilibrium concentrations required to evaluate engine-out emissions such as NO and CO along with the temporal variation of the ion current. As an illustrative example, a large-parametric study over four different design variables over sixteen different drive cycles (4096 total cases) was conducted concurrently on large computing clusters. Important trends of the design parameters on the engine performance and emissions were demonstrated. Excellent weak scaling was demonstrated on two different architectures. Furthermore, it was shown that computations of a fleet of cars (range of driving conditions and drive cycles) could be completed in near real-time speeds. pMODES thus has been demonstrated to have the potential for use in the large-scale design analyses of automotive engines.

References

1. Aithal, S. M., "Development of an Integrated Design Tool for Real-Time Analyses of Performance and Emissions in Engines Powered by Alternative Fuels," SAE Technical Paper 2013-24-0134, 2013. doi:10.4271/2013-24-0134
2. Aithal, S. M. and Wild, S. M., "Development of a Fast, Robust Numerical Tool for the Design, Optimization, and Control of IC Engines," SAE Technical Paper 2013-24-0141, 2013. doi:10.4271/2013-24-0141
3. Watanabe, K., Ito, S., and Tsurushima, T., "A New Quasi-Dimensional Combustion Model Applicable to Direct Injection Gasoline Engine," SAE Technical Paper 2010-01-0544, 2010. doi:10.4271/2010-01-0544
4. Dai, W., Davis, G., Hall, M., and Matthews, R., "Diluent and Lean Mixture Combustion Modeling for SI Engines with a Quasi-Dimensional Model," SAE Technical Paper 952382, 1995.
5. Jung, D., and Assanis, D., "Multi-Zone DI Diesel Spray Combustion Model for Cycle Simulation Studies of Engine Performance and Emissions," SAE Technical Paper 2001-01-1246, 2001.
6. Andersson, M., Johansson, B., Hultqvist, A., and Nöhre, C., "A Real Time NOx Model for Conventional and Partially Premixed Diesel Combustion," SAE Technical Paper 2006-01-0195, 2006.
7. Abu-Nada, E., Al-Hinti, I., Akash, B., and Al-Sarkhi, A., "Thermodynamic Analysis of Spark-Ignition Engine Using a Gas Mixture Model for the Working Fluid," Int. J. Energy Research 31(11):1031-1046, 2007.
8. Perini, F., Paltrinieri, F., and Mattarelli, M., "A Quasi-Dimensional Combustion Model for Performance and Emissions of SI Engines Running on Hydrogen-Methane Blends," Int. J. Hydrogen Energy 35(10):4687-4701, 2010.
9. Grill, M., Bargende, M., Rether, D., and Schmid, A., "Quasi-dimensional and Empirical Modeling of Compression-Ignition Engine Combustion and Emissions," SAE Technical Paper 2010-01-0151, 2010.

10. Grill, M., and Bargende, M., "The Development of an Highly Modular Designed Zero-Dimensional Engine Process Calculation Code," SAE Int. J. Engines 3(1):1-11, 2010.
11. Afshari, M., Daryan, J. H., Jazayeri, S. A., Ebrahimi, R., & Karan, F. S. N., "A Numerical Investigation on a Spark Ignition Engine Fueled with the Hydrogen-Methane Blend Using a Quasi-Dimensional Method," SAE Technical Paper 2015-01-0770, 2015.
12. Lamb, D., Gorsich, D., Krayterman, D., Choi, K. et al., "Predicting Military Ground Vehicle Reliability using High Performance Computing," SAE Technical Paper 2007-01-1421, 2007. doi:10.4271/2007-01-1421
13. Lamb, D., Gorsich, D., Krayterman, D., Choi, K. et al., "System Level RBDO for Military Ground Vehicles using High Performance Computing," SAE Technical Paper 2008-01-0543, 2008. doi:10.4271/2008-01-0543
14. Aithal, S. M., "Analysis of The Current Signature in a Constant-Volume Combustion Chamber," Combustion Science and Technology 185(2):336-349, 2013.
15. Aithal, S. M., "Prediction of Voltage Signature in a Homogeneous Charge Compression Ignition (HCCI) Engine Fueled with Propane and Acetylene," Combustion Science and Technology 185(8):1184-1201, 2013.
16. Kong, H., "Application of Glow Plugs for Combustion Sensing in a Diesel Engine," SAE Technical Paper 911878, 1991.
17. Morris, J., "Intra-Cylinder Combustion Pressure Sensing," SAE Technical Paper 870816, 1987. doi:10.4271/870816
18. Yoshiyama, S., "Detection of Combustion Quality in a Production SI Engine Using Ion Sensor," SAE Technical Paper 2010-01-2255, 2010.
19. Nutton, D. and Pinnock, R., "Closed Loop Ignition and Fueling Control Using Optical Combustion Sensors," SAE Technical Paper 900486, 1990.
20. Moawad, A., Balaprakash, P., Rousseau, A., and Wild, S. M., "Novel Large-Scale Simulation Process to Support DOT's CAFE Modeling System," Proc. Int. Electric Vehicle Symp. & Exhib., May 2015.
21. Homem-de-Mello, T., and Bayraksan, G., "Monte Carlo sampling-based methods for stochastic optimization," Surv. Oper. Res. & Man. Sci. 19:56-85, 2014.
22. Ehrgott, M., "Multicriteria Optimization", Springer-Verlag, 2nd ed., 2005.
23. Vijayagopal, R., Sharer, P., Wild, S. M., Rousseau, A., Chen, R., Bhide, S., Dongarkar, G., Zhang, M., and Meier, R., "Using Multi-Objective Optimization for HEV Component Sizing," Proc. Int. Electric Vehicle Symp. & Exhib., no. EVS28 0153, May 2015.

Acknowledgments

This publication was created (or partially created) by UChicago Argonne, LLC as manager of Argonne National Laboratory for the U. S. Department of Energy, Office of Science, under Contract No. DE-AC02-06CH11357. We gratefully acknowledge the computing resources provided by the Argonne Leadership Computing Facility and the Laboratory Computing Resource Center at Argonne National Laboratory.